

DOCKET NO. CH9-2000-0004 (246)

IMPROVED SPEECH RECOGNITION BY AUTOMATED CONTEXT CREATION

Inventors:
Dieter Jaepel
Juergen Klenk

International Business Machines Corporation

IBM DOCKET NO. CH9-2000-0004
IBM DISCLOSURE NO. CH8-1999-0096

Express Mail Label No.: EL920516495US

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of European Application No. 00116450.8, filed July 28, 2000 at the European Patent Office.

BACKGROUND OF THE INVENTION

Technical Field

The present invention relates to the field of speech processing and speech recognition in general. In particular, the invention relates to systems and methods for generating an output by means of a speech input.

Description of the Related Art

Due to recent advances in computer technology, as well as recent advances in the development of algorithms for speech recognition and processing, speech recognition systems have become increasingly more powerful while becoming less expensive. Certain speech recognition systems can match the words to be recognized with words of a vocabulary. The words in the vocabulary usually are represented by word models, which can be referred to as word baseforms. For example, a word can be represented by a sequence of Markov models. The word models can be used in connection with the speech input in order to match the input to the words in the vocabulary.

Most of today's speech recognition systems are continuously being improved by providing larger vocabularies or by increasing the recognition rate by employing improved algorithms. Such systems typically can include 100,000 words. Other products, for example the ViaVoice family of software available from International Business Machines Corporation, can include approximately 240,000 word entries. Many commercially available speech recognition systems operate by comparing a spoken utterance against each word in the system's vocabulary. Since each such comparison can require thousands of computer instructions, the amount of computation

required to recognize an utterance grows dramatically with increasing vocabulary size. This increase in computation has been a major problem in the development of large vocabulary systems.

Some speech recognition systems can be trained by the user uttering a training text of known words. Through this training process, the speech recognition system can be tailored to a particular user. Such training can lead to an improved recognition rate. Additionally, there are bi-gram and tri-gram based recognition systems that can search for like-sounding words such as 'to', 'two', and 'too', by analyzing such words in a context of two consecutive words (di-gram technology) or three consecutive words (tri-gram technology). The di-gram technology and the tri-gram technology also can lead to an improved recognition rate.

One problem of conventional speech recognition systems can be that as the system vocabulary grows, the number of words that are similar in sound also tends to grow. As a result, there is an increased likelihood that an utterance corresponding to a given word from the vocabulary will be mis-recognized as corresponding to another similar sounding word from the vocabulary.

Different approaches are known in the art for reducing the likelihood of word confusion. One such method is called "pruning". Pruning is a common computer technique used to reduce a computation. Generally speaking, pruning reduces the number of cases which are considered by eliminating some cases from further consideration. Scores (representing the likelihood of occurrence in an input) can be assigned to the words in a vocabulary. The scores can be used to eliminate words from consideration during the recognition task. The score can be updated during the recognition task and words which are deemed irrelevant for the recognition are not considered any further.

Another technique used to cope with large vocabulary systems is that of hypothesis and test, which is, in effect, also a type of pruning. When features are observed in a speech input, the features are used to form a hypothesis that the word

actually spoken corresponds to a subset of words from the original vocabulary. The speech input can be processed further by performing a more lengthy match of each word in this sub-vocabulary against the received acoustic signal. This sub-vocabulary is directly derived from the speech input.

5 Yet another approach for dealing with the large computational demands of speech recognition in large vocabulary systems, is the development of special purpose hardware to increase significantly the speed of such processing. There are for example special purpose processors that perform probabilistic frame matching at high speed.

10 There are a host of other problems which have been encountered in known speech recognition systems. These problems can include, but are not limited to, background noise, speaker-dependent utterance of words, and insufficient processing speed. All of these disadvantages and problems have so far prevented widespread use of speech recognition in many market domains. Accordingly, despite the recent
15 advances in speech recognition technology, there is a great need to improve further the performance of speech recognition systems before such systems find larger distribution in the market.

SUMMARY OF THE INVENTION

It is an object of the present invention to provide a speech processing system and method having an increased ease of use. The method according to an illustrative embodiment of the present invention provides a procedure where a voice-generated output can be generated using a computer system. The output can be generated by receiving an input and automatically creating a context-enhanced database using information derived from the input. The voice-generated output can be generated from a speech signal by performing a speech recognition task to convert the speech signal into computer processable segments. During this speech recognition task, the context-enhanced database can be accessed to improve the speech recognition rate. For example, the speech signal can be interpreted with respect to the words included within the context-enhanced database. Additionally, a user can edit or correct the output to generate a final output which can be made available.

A speech processing system, in accordance with the present invention, can produce a voice-generated output. The system can include a module for automatically creating a context-enhanced database by using information derived from a system input. A speech recognition system for converting a speech signal into segments also can be included. The context-enhanced database can be accessed to find matching segments. The system further can include a module for preparing the voice-generated output with the matching segments and a module for enabling editing and/or correction of the output to generate a final output. The final output, or speech-generated output can be made available.

According to the present invention, the number of words which undergo an extensive match, for example an acoustic match, against uttered words can be drastically reduced. Using the present invention, speech recognition system implementations can be provided that are less expensive and computationally less demanding. In other words, the present invention can be used in smaller systems which are less powerful than presently available desktop computers. Advantages of the

DOCKET NO. CH9-2000-0004 (246)

present invention are addressed in connection with the detailed description or are apparent from the description.

09/10/07 07:20:04

BRIEF DESCRIPTION OF THE DRAWINGS

There are shown in the drawings, embodiments which are presently preferred, it being understood, however, that the invention is not so limited to the precise arrangements and instrumentalities shown.

5 Figure 1 shows a schematic block diagram of a conventional speech recognition system.

Figure 2 shows a schematic block diagram of a first speech processing system according to the present invention.

10 Figure 3 shows a schematic block diagram of a second speech processing system according to the present invention.

Figure 4 shows a schematic block diagram of a building block of the first speech processing system according to the present invention.

Figure 5 shows a schematic block diagram of a third speech processing system according to the present invention.

15 Figure 6 shows a schematic block diagram of a fourth speech processing system according to the present invention.

Figure 7 shows a schematic block diagram of a fifth speech processing system according to the present invention.

20 Figure 8 shows an exemplary graphical user interface (GUI) for use with a speech processing system according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

According to the present invention, a scheme is provided that greatly simplifies the interaction between a user and a computer system. It is herein proposed to use available input information to provide improved and more accurate speech recognition. If a user works with a computer system, usually, there is at least one active application program, i.e., a program that is currently being used by the user. It is assumed that the user is working on or with this active application program. In many cases, the active application program can be closely related to the user's current work task. This can be illustrated by means of a simple example. Assuming that the user of a computer system (recipient) has received an electronic mail (E-mail) from another user, it is likely that the recipient will open the E-mail in order to print or read it. It is further assumed that the other user is expecting the recipient to respond to this E-mail. This means that the respective mailer software (e.g., Lotus Notes) can be active and that the E-mail is displayed in a window on the computer screen. It is highly likely, that the contents of this E-mail define the context for the recipient's response. Input information thus can be derived from this E-mail.

According to the present invention, input information can be derived in a pre-processing step which defines the contents for an output that is to be generated by the user of the computer system. In the above example, the input information can be derived from the text contained in the E-mail received. It is, however, also possible:

1. to derive the input information from a history of E-mails (e.g., a chain of inquiries and responses to these inquiries);
2. to derive the input information from a document that is currently on the computer screen (e.g., a scientific paper currently read by the user);
3. to derive the input information from a chain of related documents;
4. to derive the input information from linked documents;
5. to derive the input information from a specific folder or directory;
6. to derive the input information from attachments that are received with an E-mail;

7. to derive the input information from a spread sheet currently used by the user;
8. to derive the input information from the computer cache memory;
9. to derive the input information from the history information recorded by a web browser;
10. to derive the input information from a knowledge management system;
11. to derive the input information from an incoming message, e.g., an incoming request in a call center;
12. to derive the input information from a received facsimile;
13. to derive the input information from the result of a database search; and so forth.

This input information, no matter how it is generated, is assumed to define the context in which the user is expected to generate an output as mentioned above. According to the present invention, the user is enabled to generate this output by uttering words. The respective output thus can be referred to as a voice-generated output. For example, the voice-generated output can be an E-mail, a facsimile, a letter, a memo, or any other output (e.g., a reaction) that can be generated by a computer system.

To prepare the voice-generated output, the user is requested to utter words. This speech input undergoes a speech recognition task after having been transformed from an audio signal into a signal format that can be processed by a computer system. For this purpose, an audio system is employed. The audio system can include a microphone, a microphone followed by some audio processing unit(s), or similar means. The audio system is employed by the speech recognition system to receive the words uttered by the user, to transform the uttered words into an audio signal, and to feed this audio signal to a converter system.

The converter system can include an analog-to-digital (A/D) circuit, for example. The converter system can convert the audio signal into a signal format that can be processed by the computer system. In most implementations of the speech recognition

system, and according to the present invention, the converter system generates a digital signal.

According to the present invention, a speech recognition task can be performed to convert the uttered words into computer-processable segments, such as word segments (e.g., letters or syllables), phonemes, phonetic baseforms, frames, nodes, frequency spectra, baseforms, word templates, words, partial sentences, and so forth. Computer-processable in the present context means a representation that can be processed by a computer system.

To perform speech recognition tasks in an efficient and reliable manner, a context-enhanced database can be generated using the input information received. The context-enhanced database can be directly derived from the input information, or can be derived from an existing database using the input information. The input information can be used, for example, to define a smaller, specific portion within a pre-installed larger lexicon. A context-enhanced database can include a few words up to several thousand words, preferably between 10 words and 1,000 words. The size of the context-enhanced database can depend upon the actual implementation of the inventive scheme and on the size of the input itself. According to the present invention, the context-enhanced database can be dynamically generated or updated depending on, or taking into account, the user's current or most recent activities.

As previously mentioned, the context-enhanced database can be generated directly from the input information or can be derived from an existing database using the input information. The latter can be done by generating a word list from the input information (e.g., by extracting words from an E-mail to be responded to) and by connecting or linking this word list to an existing lexicon. The word list can be connected or linked to the lexicon such that it acts as a filter or a first instance that can be accessed during a speech recognition task. In that case, the underlying lexicon need only be accessed if no matching word was found in the word list. Other ways of implementing this aspect of the invention will be discussed in further detail below.

During the speech recognition task, the context-enhanced database can be accessed in order to improve the speech recognition rate. The segments derived from the words uttered by the user when preparing an output can be interpreted in light of the words given in the context-enhanced database. According to the present invention, the number of processable segments which undergo an extensive match (e.g., an acoustic match) against uttered segments can be drastically reduced, since the matching is done - at least in a first run - with information in the context-enhanced database only.

According to the present scheme, the output can be prepared while the user talks into the audio system. In a subsequent step, the system can enable the user to edit or correct the output in order to generate a final output. There are different approaches that can be used to enable a user to edit or correct the output. The system can, for example, display the output on a screen to allow the user to read it and to intervene manually if there is something to be edited or corrected. Also, the system can highlight those words where there is a certain likelihood of misinterpretation (mis-recognition) of the user's speech, for example in the case of unknown words, similar sounding words, and the like. Other implementation examples are given in connection with specific embodiments.

After having finished the speech recognition task, the final output is made available for further processing. The final output can be sent via a mailer to another user, prepared for printing, mailed via a fax modem or a fax machine, stored in a memory, and so on. For this purpose, the output can be temporarily put into a memory from where it can be printed, transmitted, fetched by some other application program, or the like.

The present invention can improve known speech recognition schemes by providing a context-enhanced database which is derived from some input information that is assumed to be related to the user's current task. Thus, speech recognition can be performed in light of a well defined context rather than a huge lexicon. An output is

generated by transcribing or synthesizing the human dictation in light of the a context-enhanced database. The expression "computer system" as used herein can be a synonym for any system that has some computational capability. Examples can include, but are not limited to, personal computers (PCs), notebook computers, laptop computers, personal digital assistants (PDAs), cellular phones, and the like.

A speech recognition system is a system that performs a speech recognition task. Typically, a speech recognition system is a combination of a general purpose computer system with speech recognition software. A speech recognition system also can be a special purpose computer system, such as a system with special purpose speech recognition hardware.

Speech recognition systems are marketed which can run on a commercial PC and which require little extra hardware except for an inexpensive audio system, for example a microphone, an audio card with an analog-to-digital (A/D) converter, and a relatively inexpensive microprocessor to perform simple signal processing tasks. Such systems can provide discrete word recognition. There are also computer systems which require just speech recognition software. The necessary hardware components are already present in the form of an integrated microphone and an A/D converter.

A schematic representation of a conventional speech recognition system 10 is illustrated in Figure 1. Most speech recognition systems 10 operate by matching an acoustic description of words (e.g. Word 14) in their lexicon 13 against a representation of the acoustic signal generated by the utterance of the word (e.g. Word' 11 received as an input 12) to be recognized. If the input word 11 matches the word 14 in the lexicon 13, then an output 15 can be generated including the word 14. The speech signal representing the word 11 can be converted by an A/D converter into a digital (signal) representation of the successive amplitudes of the audio signal created by the speech. That digital signal can be converted into a frequency domain signal composed, at least in part, of a sequence of frames, each of which gives the amplitude of the speech signal in each of a plurality of frequency bands. Such systems commonly operate by

comparing the sequences of frames produced by the utterance to be recognized with a sequence of nodes, or frame models, contained in the acoustic model of each word in their lexicon. Such a speech recognition system is called a frame matching system.

The performance of frame matching systems can be improved using a probabilistic matching scheme and a dynamic programming scheme, both of which have been known in the art for some time now. The application of dynamic programming to speech recognition is described in the article "Speech Recognition by Machine: A Review" D.R. Reddy, in Readings in Speech Recognition, A. Waibel and K.-F. Lee, Editors, 1990, Morgan Kaufmann: San Mateo, CA, pp. 8 - 38.

One embodiment of a speech processing system 20, according to the present invention, is illustrated in Figure 2. In a pre-processing step, a context-enhanced database 21 can be generated from input information 22, as described in one of the previous sections. If now a speech signal is received by an audio system 24, as indicated by arrow 23, an audio signal representing this speech signal can be forwarded via line 25 to a converter system 26. The audio system 24 can transform the acoustic signal received via 23 into the audio signal. This audio signal can be fed via line 25 to the converter system 26 where it is transformed into a signal format that is processable by the speech recognition engine 27. In most implementations, the converter system 26 is designed to generate a digital signal that is fed via line 28 to the speech recognition engine 27. This digital signal represents processable segments uttered by a user.

The speech recognition engine 27 can match the processable segments with segments in the context-enhanced database 21, as indicated by the arrow 29. All those segments for which a matching segment was found in the context-enhanced database 21 (called matching segments) can be fed to an output unit 30 where an output is generated. The user now can interact with the system 20 by editing and/or correcting the output, as indicated by the output editing/correction unit 31. The user interaction is illustrated by the arrow 32. The unit 31 can provide a final output 33 at an output line

34. Depending on the implementation, some of the steps can be performed concurrently.

Another embodiment of a speech processing system 40, according to the present invention, is illustrated in Figure 3. In this example, an E-mail 41 from an E-mail folder 42 delivers the input information. The E-mail folder 42 can be the inbox of a mailer 43. The mailer 43 also can have an outbox 44, which can include at least one E-mail 45 waiting for delivery. As schematically shown in Figure 3, the E-mail 41 contains the usual address information 47, a subject field 46, and a text body 48. According to the present embodiment, a word list 49 can be derived from the contents of the E-mail 41. The word list 49 can be derived from the address information 47, the subject field 46, the text body 48, or from any combination thereof. This word list 49 can be used to provide a context-enhanced database (not shown). In the present embodiment, the word list 49 sits on top of a lexicon 13 that has many word entries.

If the user now wants to prepare an output (e.g., a response to the E-mail 41), the user can, for example, activate the speech recognition module and talk into a microphone. The respective speech signal (box 50) can be analyzed by a conventional phoneme processing engine 51. Then a word matching process can be carried out by the word matching engine 52. This word matching engine 52 can include an application programming interface (API) 53 that serves as an interface for accessing a lexicon. A conventional speech recognition system can access a large lexicon, for example lexicon 13, through the interface 53 to find matching words. According to the present invention, however, the word list 49 can be accessed first through the API interface 53. If all words uttered by the user and represented by the speech signal are found in the word list 49, a grammar check can be performed by a grammar check unit 54 before an output 57 can be generated by the output generation unit 55. This output 57 can be provided at the output line 56 for further processing. In the present embodiment, the output 57 is the body of an E-mail that is stored in a memory unit 58. It can be fetched from this memory 58 and pasted into an outgoing E-mail. The E-mail 45 that sits in the

outbox of the mailer 43 was generated exactly the same way. As soon as the computer system 40 connects to a network, the outgoing mail can be transmitted.

The word matching engine 52 can be implemented such that it always returns the best match for a word received from the unit 51. Part of the output can be presented to the user right away at output line 56 before the user has completed spelling the desired words. Advantageously, the speech processing system 40 can be implemented in such a way that the lexicon 13 can be accessed if there are words for which no matching counterpart was found in the word list 49. This can be done through the same API interface 53, or a separate API interface which can be provided for that purpose.

The pre-processing module 36, which performs the pre-processing steps described in connection with the embodiment of Figure 2, is schematically summarized in Figure 4. As illustrated in Figure 4, some input information 22 can be received via an input line 35. This input information 22 can stem from an E-mail currently processed in an editor or from a some other source, as indicated by the aforementioned items 1 - 13. The context-enhanced database 21 can be automatically created by deriving information from the input information 22. There can be an interface 29 which allows the speech recognition engine 27 (cf. Figure 2) to access the context-enhanced database 21. This context-enhanced database 21 can return matching segments or matching words.

Another pre-processing module 65 is shown in Figure 5. The input information 22 is received via an input line 63. A processing unit 60 can be employed which can take information (e.g., segments or words) from the input information 22 and create a context-enhanced database 62. In order to obtain an improved context-enhanced database 62, a synonym lexicon 61 can be employed. If the input information comprises a Word A, the processing unit 60 can create several entries in the context-enhanced database 62; one for the original word, namely Word A, and as many entries as there are synonyms in the synonym lexicon 61. Assuming that there are three

synonyms Word A', Word A'', and Word A''' in lexicon 61, four entries (Word A, Word A', Word A'', and Word A''') can be created in the context-enhanced database 62. In accordance with the present invention, if the user of a system speaks a word that is a synonym to a word included in the input information 22, the system can recognize this word and add it to the output being generated. Interface 64 (e.g. a standardized interface) can allow a speech processing system to access the context-enhanced database 62.

Yet another pre-processing module 75 is depicted in Figure 6. In this example, the input information 22 is received via an input line 73. A processing unit 70 can be employed which can take information (e.g., segments or words) from the input information 22 and build a context-enhanced database 72. In order to obtain a context-enhanced database 72 with more word entries, a database 76 with meaning variants and a synonym lexicon 77 can be employed. If the input information includes a Word A (e.g., the word "plant"), the processing unit 70 can access the meaning variants database 76 in order to check whether there is more than one meaning for the Word A. In case of the word "plant", for example, the database 76 can include two entries. The first entry (Word A*) can identify the "living plant" and the second entry (Word A**) can identify the "building" or "industrial fabrication plant". Both of these meaning variants (Word A* and Word A**) can be retrieved by the processing unit 60. Other information can be used by the processing unit 60 to identify which of the two variants (Word A* or Word A**) is the one that is actually meant. If the input information contains the sentence "A plant was erected in 1985", for example, then it is clear from the context that the building (Word A**) and not the living object (Word A*) is referenced. The synonym lexicon 77 now delivers synonyms for this second variant (Word A**). This scheme allows the system to avoid misunderstandings due to the misinterpretation of different word variants. It resolves these issues while creating the context-enhanced database 72. The context-enhanced database 72 can be accessible via the interface 74.

The embodiment illustrated in Figure 7 can be more powerful than the previous embodiments, since it employs a meaning extraction system 81 in connection with a knowledge database 86. The module 85 can include a memory with the input information 22. The processing unit 80 can consult the meaning extraction system in order to get some understanding of what is contained in the input information 22. The meaning extraction system can be a system interacting with a fractal hierarchical knowledge database 86, as for example described and claimed in the European patent application entitled "Processing of textual information and automated apprehension of information", filed on June 2, 1998, and which is currently assigned to the assignee of the instant patent application. Such a meaning extraction system 81 can understand - at least to some extent - what is meant by the input information 22. It further can extract additional information believed to be associated or related. Thus, the processing unit 80 can build a context-enhanced database 82 that is 'richer' in that it not only contains the words that were found in the input information, but also information that is deemed to be related. The context-enhanced database 82 can be accessible via the interface 84. The input information 22 can be received via an input line 83.

An example of a graphical user interface (GUI) for use with a simple speech recognition system is illustrated in Figure 8. The GUI can include an editor 90 with a text-editing window 91. The speech recognition system can display the result of a speech recognition exercise where the user has uttered the partial sentence "...plant causes pollution ...".

As shown in Figure 8, recognition of the user uttered sentence was, for the most part, accurate, with the exception of the word "plant", which was misunderstood by the system. As shown in the text-editing window 91, the system can transform the partial sentence and can display the resulting output 92. As shown, the system has output the word "plan" rather than the word "plant". The user now can edit/correct the output 92. To do so, the user can double-click on a word that is believed to be misspelled. The

system can highlight the respective word. In the present example, the word "plan" is highlighted using the mouse. A correction window 93 can be opened that offers different alternatives 1 - 2 for the word "plan". If a meaning extraction system, similar to one in Figure 7, is employed in the background, the processing system can tell from the context of the context-enhanced database that an industrial fabrication plant was intended. The system thus can offer the best matching word in the uppermost position 1 in the correction window 93. The system can determine that the other word "planned" is not likely to be relevant since this word does not make sense in the present context. By clicking on the OK-button, the word "plan" can be corrected so that it reads "plant".

A speech recognition system according to the present invention can be realized such that the word "plan" is automatically corrected. This can be achieved because the system can recognize that the word "plant" is the only word that makes sense in the present context.

An implementation of the present invention that makes use of a word list (context-enhanced database) generated from an active window (e.g., an E-mail) can check whether the word "plan" is included in the context-enhanced database. If this word is not in the context-enhanced database, the system can replace it with the word "plant", provided that the word "plant" is in the context-enhanced database. A system according to the one illustrated in Figure 7 can determine that the combination of the words "pollution" and "plant" is valid and that the combination of "pollution" and "plan" is not valid. This capability also allows for automatic corrections.

According to one embodiment of the present invention, a template (form) can be generated automatically from the input information. The voice-generated output can be inserted into the template. Such a template-based approach can be well suited for situations where a highly automated response is required and where the responses typically look the same. An example could be a booking system used by a chain of affiliated hotels.

The present invention can be used in connection with systems that process discrete speech (e.g., word-by-word) or continuous speech. Advantageously, a system according to the present invention can include a speech synthesizer that converts the final output into a speech output. Such a speech synthesizer can include synthesizer hardware with a parameter store containing representations of words to be output, as well as a loudspeaker, for example.

Another embodiment of the present invention can include a fall-back mode or procedure which can be engaged in those situations where no matching words are found. Such a fall-back mode or procedure can offer the user a simple interface for typing the missing words.

According to another embodiment of the present invention, the context-enhanced database can be dynamically generated while input information is received. A first guess context-enhanced database can be generated and then constantly updated as additional input information is received. For example, a call can be received on a call-in line of a call center. The call center system can route the call to an automated call handler which asks questions. The caller can respond by uttering words or alternatively by pressing buttons on the phone. While this interaction continues, a first guess of a context-enhanced database can be generated. If the caller is not calling for the first time, caller specific information can be fetched from a memory. This caller specific information can be used to generate a context-enhanced database, or an old context-enhanced database can be retrieved that was generated during a previous call of the same caller. The context-enhanced database can be constantly updated as the caller reveals additional information about the reason for calling. An output can be generated (e.g. a confirmation fax) by the operator of the system. In order to do so, the operator speaks into a microphone. The words he utters can be transformed and processed referring to the most current version of the context-enhanced database. The final output can be temporarily stored, printed, signed, and faxed to the caller's fax number.

Using the present invention, one is able to transcribe human dictation into an output, such as a letter or an E-mail. This greatly increases the speed and ease with which humans can communicate with other humans using computer-generated letters or E-mail. Additionally, using the present invention, humans can record and/or organize
5 their own words and thoughts. This can be done by storing a voice-generated output in a database, or by using the voice-generated output to update a knowledge database.

Another advantage of the present invention is that it can be used on PDA or phone-like systems which lack an adequate keyboard. With the proposed embodiments, the speed of retrieval and the recognition rate can be improved since the
10 context-enhanced database enables faster and more reliable matching.